

Amendments to the Specification

Please replace the paragraph beginning on page 19, line 25, and ending on page 20, line 8, with the following amended paragraph:

FIGS. 8-9 illustrate computer programs that perform validity filtering or quality filtering (or both) on system designs composed of component designs D_1, \dots, D_n . In FIG. 8, respective common component validity filters C_1, \dots, C_n prepare component validity sets for respective component designs D_1, \dots, D_n . The component validity sets are then filtered by partial component validity filters defined by partial component validity predicates $(V_{11}, \dots, V_{1a}), \dots, (V_{n1}, \dots, V_{nz})$, respectively. As noted previously, for any component design space for which all designs are known to be valid, validity filtering can be omitted and if all system designs are known to be valid, validity filtering can be completely omitted. The resulting partial component validity sets are combined to form component validity sets S_{11}, \dots, S_{nm} . In steps 801₁, \dots , 801_m Cartesian products of these sets form system design sets $S_{1, \dots, S_n} S_{1, \dots, S_m}$ that are combined to form a system validity set S.

Please replace the paragraph beginning on page 21, line 29, and ending on page 22, line 12, with the following amended paragraph:

Common terms in the validity function, such as $(instrSize \leq 64)$, are evaluated with reference to a component design for a single component. The corresponding common component validity filter (one of the common component validity filters C_1, \dots, C_n) evaluates the term $(instrSize \leq 64)$ based on the processor design only, without

consideration of the memory design. The terms $(\text{intLitSize} \leq 32)$ and $(\text{memLitSize} \leq 32)$ appear to qualify as common terms but do not appear in both AND expressions. Because $(\text{intLitSize} \leq 32)$ does not appear in both AND expressions, a component design that does not satisfy the term $(\text{intLitSize} \leq 32)$ can be an element of a validity set. The result of an evaluation of a validity predicate that includes a common term is TRUE (valid) only if the common term is also TRUE (valid). Consequently, component designs that do not satisfy a common term are not part of any valid system design.

Please replace the paragraph beginning on page 28, line 14, and ending on page 29, line 3, with the following amended paragraph:

Given a Pareto curve or a comprehensive Pareto set, a design can be selected programmatically to achieve a predetermined cost or time, or combination of cost and time. Using the Pareto curve (or the comprehensive Pareto set), superior designs are not overlooked. However, construction of the Pareto curve and the comprehensive Pareto set by exhaustively evaluating all possible designs is generally infeasible due to the large number of design variables available as well as the complexity of evaluating a particular design. As shown in, for example, FIGS. 4-8, a processor system or other system of interest can be divided into components and $[[a]]$ component design spaces can be quality filtered (i.e., Pareto filtered) to produce component quality sets that are component Pareto sets. Combining the component Pareto curves or sets constructs a comprehensive Pareto curve or Pareto set for the system. For example, a system design d is a composition of component designs d^1, d^2, \dots, d^n , and a set of system

designs is obtained from the Cartesian product of sets of component designs, i.e., the set of systems designs is the set of all combinations of the component designs. The program can also determine the validity of a component design or a combination of component designs, as described previously.

Please replace the paragraph beginning on page 34, line 29, and ending on page 35, line 7, with the following amended paragraph:

Each of the cache components is evaluated individually. The d-cache 1207 is evaluated as a function of cache size only, as a direct mapped cache with a line size of 16 bytes. FIG. 13 contains a Pareto ~~curve 301~~ curve 1301 for the d-cache 1207 for cache sizes of 2, 4, 8, and 16 kB. FIG. 13 also shows Pareto designs 1303, 1305, 1307, 1309 for the d-cache 207. The Pareto curve 1301 is graphed with design execution time (d-cache misses N_d) on a vertical axis 1311 and cache cost (wafer area) on a horizontal axis 1313. An approximate Pareto curve 1315 connects the Pareto designs 1303, 1305, 1307, 1309.

Please replace the paragraph beginning on page 35, line 8, with the following amended paragraph:

The line size of the i-cache 1209 is fixed at 32 bytes. The size of the i-cache 1209 ranges from 2 kB to 64 kB and associativities of 1 and 2 are considered. The costs and execution times for these combinations of size and associativity are determined based on the number of cache misses in the i-cache 1207 as a function of

cache size based on the simulated execution of the GHOSTSCRIPT application with a predetermined design of the VLIW processor 1201. FIG. 14 contains a Pareto curve 1401 for the i-cache 1209 that is plotted with respect to coordinate axes 1405, 1407 corresponding to execution time (i.e., i-cache misses N_i) and cost, respectively. FIG. 14 also shows Pareto design points 1403 as well as non-Pareto design points 1409. The Pareto curve 1401 eclipses the non-Pareto design points 1409. As discussed above, the execution time is determined as a number of i-cache misses, i.e., the number of times the VLIW processor 1201 is unable to retrieve the requested instruction directly from the i-cache 1207 while executing the ~~GHOSTSCRIPT~~ GHOSTSCRIPT application.

Please replace the paragraph beginning on page 35, line 24, and ending on page 36, line 8, with the following amended paragraph:

For both the d-cache 1207 and the i-cache 1209, the actual execution time depends on the design of the u-cache 1211. Design of the u-cache 1211 is considered next. Design variables for the u-cache 1211 considered in this design example include cache size (64 kB to 2 MB) and associativities (2 and 4). The u-cache 1211 communicates with main memory via a system bus and requires a main memory cycle time t_{main} to retrieve data from main memory. The u-cache designs considered require an access time (t_{access}) that is equivalent to 3-7 processor clock cycles to a supply not found in the i-cache 1207 or the d-cache 1209. FIG. 15 contains a component Pareto curve 1501 for the u-cache 1211 and Pareto design points 1503, 1505, 1507, 1509, 1511 that correspond to access times of 3, 4, 5, 6, and 7 processor clock cycles,

respectively. FIG. 15 also shows non-Pareto design ~~points~~ point 1513. For convenience, the Pareto curve 1501 is shown as a smooth curve connecting the Pareto design points.

Please replace the paragraph beginning on page 37, line 13, with the following amended paragraph:

As yet another example of design selection using component Pareto sets or curves to form a comprehensive Pareto set, a design for the VLIW processor system 1200 of FIG. 12 can be selected using component Pareto sets or curves for the VLIW processor 1201, the systolic array 1203, and the combined cache 1205 to prepare a combined Pareto set. As in the previous examples, the performance criteria are cost and execution time. FIG. 17 contains graphs of the component Pareto curves. In this example, VLIW processor designs are considered having various numbers of data ports for communication with the d-cache 1209. A graph 1701 of component Pareto curves for the VLIW processor 1201 includes curves 1703, 1705 that represent component Pareto curves for different numbers of d-ports. Similarly, a graph 1711 of component Pareto curves for the systolic ~~array 203~~ array 1203 includes component Pareto curves 1713, 1715 for different numbers of systolic ports 1217.

Please replace the paragraph beginning on page 37, line 28, and ending on page 38, line 14, with the following amended paragraph:

Component Pareto curves are also prepared for the i-cache 1209, d-cache 1207, and the u-cache 1211. A graph 1721 contains a component Pareto curve 1723 for the i-cache 1209 and is prepared as described above. A graph 1731 contains component Pareto curves 1733, 1735 for the d-cache 1209, the curves 1733, 1735 corresponding to different numbers of data ports 1215. While only two curves 1733, 1735 are shown, additional numbers of data ports 1215 can be considered. The execution time of the d-cache 1209 is independent of the number of data ports 1215, but cost is not. Similarly, a graph 1741 contains component Pareto curves 1743, 1745 for the ~~u-cache 211~~ u-cache 1211 corresponding to different numbers of u-cache ports 1219. The component Pareto curves corresponding to the d-cache 1209, the i-cache 1207 and the u-cache 1211 are combined to produce comprehensive Pareto curves 1751, 1753 corresponding to different numbers of data ports 1215 and u-cache ports 1219. The combined Pareto curves 1751, 1753 are component Pareto curves with respect to the processor system 1200.